

EXTRACTING LANDMARKS FOR CAR NAVIGATION SYSTEMS USING EXISTING GIS DATABASES AND LASER SCANNING

C. Brenner and B. Elias

Institute of Cartography and Geoinformatics, University of Hannover, Germany, claus.brenner@ikg.uni-hannover.de.

Commission III, WG III/3

KEY WORDS: Car Navigation, Landmarks, GIS, Laser Scanning.

ABSTRACT

Today's car navigation systems provide driving instructions in the form of maps, pictograms, and spoken language. However, they are so far not able to support *landmark-based navigation*, which is the most natural navigation concept for humans and which also plays an important role for upcoming personal navigation systems. In order to provide such a navigation, the first step is to identify appropriate landmarks – a task that seems to be rather easy at first sight but turns out to be quite pretentious considering the challenge to deliver such information for databases covering huge areas of Europe, Northern America and Japan. In this paper, we show approaches to extract landmarks from existing GIS databases. Since these databases in general do not contain information on building heights and visibility, we show how this can be derived from laser scanning data.

1 INTRODUCTION

Modern car navigation systems have been introduced in 1995 in upper class cars and are now available for practically any model. They are relatively complex and mature systems able to provide route guidance in form of digital maps, driving direction pictograms, and spoken language driving instructions (Zhao, 1997). Looking back to the first beginnings in the early 1980s, many nontrivial problems have been solved such as absolute positioning, provision of huge navigable maps, fast routing and reliable route guidance.

However, the original concept of delivering the instructions has not changed very much. Still, spoken language instructions use a relatively small set of commands (like 'turn right now'), which only refer to properties of the street network. This is not optimal, since *i*) features of the street network typically are not visible from a greater distance due to the low driver position and small observing angle, and *ii*) the most natural form of navigation for humans is the navigation by landmarks, i.e. the provision of a number of recognizable and memorizable views along the route.

Obviously, the introduction of buildings as landmarks together with corresponding spoken instructions (such as 'turn right after the tower') would be a step towards a more natural navigation. As we argue below, this would be well integrable into today's car navigation systems as it would not imply a major modification of systems and data structures. Thus, the main problem lies in identifying suitable landmarks and evaluating their usefulness for navigation instructions. In this paper, we show how existing databases can be exploited to tackle the first problem, while laser scanning data can be used to approach the second.

2 NAVIGATION USING LANDMARKS

There are two different kinds of route directions to convey the navigational information to the user: either in terms of a description (verbal instructions) or by means of a depiction (route map). According to (Tversky and Lee, 1999) the structure and semantic content of both is equal, they consist of landmarks, orientation and actions. Using landmarks is important, because they serve

multiple purposes in wayfinding: they help to organize space, because they are reference points in the environment and they support the navigation by identifying choice points, where a navigational decision has to be made (Golledge, 1999). Accordingly, the term landmark stands for a salient object in the environment that aids the user in navigating and understanding the space (Sorrows and Hirtle, 1999). In general, an indicator of landmarks can be particular visual characteristic, unique purpose or meaning, or central or prominent location.

Furthermore, landmarks can be divided into three categories: visual, cognitive and structural landmarks. The more of these categories apply for the particular object, the more it qualifies as a landmark (Sorrows and Hirtle, 1999). This concept is used by (Raubal and Winter, 2002) to provide measures to specify formally the landmark saliency of buildings: the strength or attractiveness of landmarks is determined by the components visual attraction (e.g. consisting of façade area, shape, color, visibility), semantic attraction (cultural and historical importance, explicit marks, e.g. shop signs) and structural attraction (nodes (important intersection), boundaries (parting elements like rail tracks or rivers), regions (building blocks)). The combination of the property values leads to a numerical estimation of the landmark's saliency.

A study of (Lovelace et al., 1999) includes an exploration of the kinds and locations of landmarks used in instructions. It can be distinguished between four groups: choice point landmarks (at decision points), potential choice point landmarks (at traversing intersections), on-route landmarks (along a path with no choice) and off-route landmarks (distant but visible from the route). A major outcome of the study is that choice point and on-route landmarks are the most used ones in route directions of unfamiliar environments.

The choice of an appropriate landmark depends on the navigation context and application mode: pedestrians or car drivers. Accordingly, there are different studies for both user groups, dealing with the when, why and how landmarks are used in instructions. Because of the different conditions (moving speed, visual field, arbitrary movement or constrained to road network), studies targeted at pedestrians (Michon and Denis, 2001, Lovelace et al.,

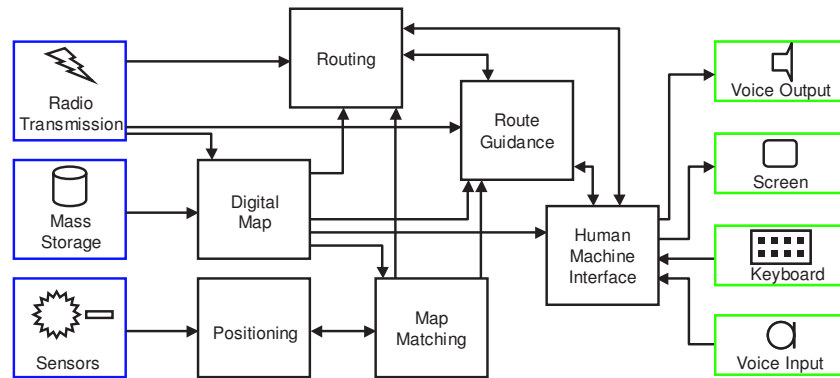


Figure 1: Block diagram of a modern car navigation system.

1999, Winter, 2002) as well as car drivers (Burnett, 1998, Burnett et al., 2001) have been undertaken. The study of (Burnett, 1998) reveals some of the underlying factors for 'good' landmarks which should be considered for designing route guidance systems. Some of the important factors are permanence, uniqueness and visibility of the landmark.

3 CAR NAVIGATION SYSTEMS

3.1 Components of Car Navigation Systems

Figure 1 shows the components of a modern car navigation system. Typically, the car's position is determined by combining signals from a GPS receiver, an angular rate sensor, as well as an odometer (speed) signal from the wheels. Since the absolute position given by GPS might be quite wrong, especially in densely buildup areas, it is corrected so as to fit to the digital map which is nowadays usually obtained from an onboard mass storage such as a CD or DVD. This process is called map matching and is realized as a multipath matching which always tracks (and rates) several possible positions in the map simultaneously. Altogether, this positioning is most of the time sufficiently accurate and reliable, working even during longer GPS outages.

One important module in Fig. 1 is the *route guidance module*. It is given an ordered list of edges to be driven from the *routing (route planning) module* as well as the current position from the positioning / map matching module. From that, it decides when to issue which instructions to the driver. For natural language instructions, this can be divided into "early warnings" (such as 'keep right' or 'prepare to turn right') and "immediate instructions" (such as 'turn right now'). When thinking about landmark-based car navigation, one has to keep in mind that this functionality has to be provided as well. Particularly, landmarks are only useful for driving instructions such as turns if they are visible at a sufficient distance from the decision point, and they do not disappear as the driver is approaching this point.

3.2 Digital Maps

The maps used by car navigation systems not only contain the geometry and connectivity of the road network but also a huge amount of additional information on objects, attributes and relationships. A good overview can be obtained from the European standard *GDF*, see e.g. (Geographic Data Files 3.0, 1995). Of particular interest are points of interest (POI) which include museums, theaters, cultural centers, city halls, etc.

Map data is acquired by map database vendors such as Tele Atlas or NavTech and supplied to car navigation manufacturers in an

exchange format (such as GDF). There, it is converted to the proprietary formats finally found on the map CD or DVD. This conversion is highly nontrivial since the data has to be transformed from a descriptive form into a specialized form supporting efficient queries by the car navigation system. Often, structures and values are precomputed by this conversion process in order to relieve the navigation system's online resources such as bandwidth and CPU time.

Part of this process is also to generate a matrix for *each intersection* which describes all possible turn combinations. Also, for the well-known arrow pictograms used by car navigation systems, the angles between all streets joining at an intersection are stored.

It is during this conversion process where additional information for landmark-based navigation can be integrated. In this paper, we outline how the street geometry given by GDF can be combined with information from a cadastral map and laser scan data to identify suitable landmarks. An important point is that the additional datasets are used only during the conversion process. After that, only landmark-based driving instructions remain, which can be coded in a very compact form and are compatible with the per-intersection information already stored in proprietary map formats. Thus, the technical integration of landmark-based instructions into current car navigation systems poses no major obstacles, and the main problem is to derive those instructions in some automatic or at least semiautomatic way.

4 LASER SCANNING AND CITY MODELS

During the 1990's, *airborne laser scanning* became available as a new method for obtaining surface models. Subsequently, the scanning systems were improved and direct georeferencing became feasible with sufficient accuracy. Today, airborne laser scanning is a mature technology with a multitude of companies offering systems and services (Baltsavias, 1999). Scanning of very large areas is possible, for example the entire Netherlands have been and Germany's state of Baden-Württemberg is in the progress of being scanned, each with an area of over 30.000 km². Aerial laser scanners produce dense point clouds of the earth's surface directly (Baltsavias et al., 1999). They are particularly suitable for obtaining digital surface models (DSMs) in dense urban areas, as they conserve jump edges quite well. Most systems are capable of measuring not only the height, but also the reflectance, as well as first, last or multiple return pulses, which allows to separate tree canopy and ground (Kraus and Rieger, 1999).

The main problem is how to extract symbolic information about man-made structures from laser scanner datasets, possibly combined with aerial or terrestrial images. Especially, the automatic

generation of city models has been and still is an intense research field, the discussion of which is beyond the scope of this paper. The reader is referred to the excellent proceedings of the “Ascona workshops” on this topic (Grün et al., 1995, Grün et al., 1997, Baltsavias et al., 2001).

However, there is still substantial research effort necessary until highly automated object extraction systems working reliably become available. On the other hand, three-dimensional object information is still far from being common in today’s existing GIS databases. In consequence, in this paper we consider using two-dimensional GIS databases in combination with laser scanner DSMs on an iconic level, without explicitly reconstructing the three-dimensional shape of the objects as separate entities. Figure 2 shows an example of the data sources used, which is a DSM from laser scanning, regularized to a 1 m grid, the street geometry represented by center lines from a GDF data set, and the outline of buildings from a cadastral map.



Figure 2: Laser scan (regularized to 1 m rasterwidth, shown color-coded), streets from GDF dataset (white) and building groundplans from cadastral map (black). Image shows part of Stuttgart, Germany.

5 LANDMARKS DISCOVERY FROM CADASTRAL MAPS

Most of the research cited in section 2 deals with theoretical aspects of navigation using landmarks. Another point of view is to investigate existing GIS databases and to extract objects that are potential landmarks, because they match the stated requirements: being salient in their environment.

A first study on the use of the topographic data base ATKIS (AdV, 2003) from the German national mapping agency was presented in (Elias, 2002). The findings point out the possibilities to use the content of ATKIS to enrich wayfinding instructions with landmarks. Here, we want to focus on building data which is part of the cadastral map and can be combined with laser scanning data.

The goal is to detect automatically all buildings in the data base which fulfill the criteria for being a landmark. Therefore, data mining methods are used to analyze the data using the geometric, topologic and the semantic information given in the data base. The results provide ‘potential’ landmarks computed by relative uniqueness of objects in their neighborhood. This calculation is completely independent from the chosen route, the situation and the way of moving. The second step is the selection of the appropriate landmark according to the current situation (route, visibility, distance to decision point, etc.) and results in the ‘real’ landmark referred to in a route guidance information.

Here, we will introduce our approach to extract landmarks, the chosen cadastral data base and their information content, as well as a description of the data mining methods used.

5.1 Approach for Extracting Landmarks

The term *knowledge discovery in databases* (KDD) can be defined as the discovery of interesting, implicit, and previously unknown knowledge from large databases (Frawley et al., 1991). It comprises the overall process of finding and interpreting patterns from data, while *data mining* only refers to the stage of data analysis without the additional steps. So the KDD process includes preprocessing of data, data mining itself and postprocessing, as well as the interpretation of potentially discovered patterns.

Data mining methods are algorithms designed to analyze data or to extract from data patterns into specific categories (Fayyad et al., 1996). Basic models of data mining are clustering, regression models, classification, summarization, link and sequence analysis. The algorithms can be divided into two basic techniques. According to the terminology of the machine learning community, there are methods for learning from examples (*supervised learning*) and learning from observation (*unsupervised learning*). In our approach, we try to extract the landmarks with unsupervised learning methods using a modified application of the *ID3* decision tree algorithm (primarily a supervised method) (Quinlan, 1986) and the clustering algorithm *Cobweb* (Witten and Eibe, 1999).

The idea is that objects, which have a unique attribute in a certain environment, qualify as landmarks. Therefore, the underlying model is to compare the attribute values of data records: Objects with distinct or even unique values in a certain spatial neighborhood are assumed to be prominent. The procedure will also lead to an attribute ranking according to their importance for the model. If the chosen attributes are suitable for developing an object schema and outliers from this schema present something particular, it is possible to determine landmarks through statistics and data mining methods. In our case, we concentrate on building objects. We enrich the explicitly given information about buildings by deriving attributes and relations with the help of spatial analysis. The combination of different attribute-values leads to derived attributes, e.g. building length to width ratio.

5.2 Original and Derived Contents of Cadastral Maps

The digital cadastral map of Lower Saxony is an object oriented vector database of state-wide availability. This digital map includes buildings, parcels and land use. Besides geometry information, the following semantic attributes are available:

- building use: residential, public, underground, outbuildings
- land use types: public purposes, residential, commerce and service, industrial, mixed land uses, traffic, park, garden, sports, etc.
- building labels: name or function of building (e.g. town hall, kindergarten, church)
- special building parts with a roof: winter garden, car port etc.

To extend the content of the data base and provide more attributes for the data mining process, the following implicitly contained information is extracted from the digital cadastral map: information




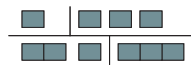



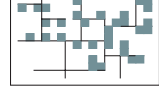

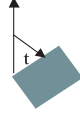
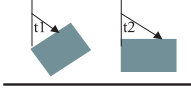



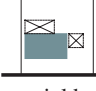

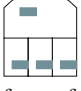

 size of building	 building form	 corners	 (semi-)detached	 road distance	 ratio (building/parcel)
 neighborhood density	 district density	 orientation to road	 orientation to north	 orientation to neighbor	 perpendicular angles
 parcel land use	 number of buildings	 special buildings	 neighbor land use	 form of parcel	 church name or function

Figure 3: Derived attributes and relations of buildings

about the geometry, semantics and topology, in this case neighborhood relations to other buildings and other object groups (for example, the distances to roads and parcel boundaries) are collected. The derived attributes are shown in Figure 3.

One profound disadvantage of the cadastral database content is the lack of height values for the buildings. This information can be provided by laser scanning data sets.

5.3 Data Mining

The prepared data is processed using the classification method ID3 and a clustering approach called Cobweb based on the data mining software WEKA (Witten and Eibe, 1999).

ID3 is a method for supervised learning and therefore needs classified examples. The result is a decision tree providing the shortest optimal description possible for a classification into the given classes. As the used data have no classified examples, we use this procedure iteratively and enhance our data set with the class landmark (values: yes, no). Therefore, we iteratively hypothesize each building to be a landmark, whereas all the other buildings are no landmarks. 'Real' / true landmarks then are identified by yielding the most simple / shortest description. This method needs a lot of computing depending on the number of buildings. The application of this technique was tested on a synthetic data set. The conclusion of the study has been that a landmark is characterized by a short decision tree with only a few levels that leads to positive landmark decision.

Cobweb is a hierarchical clustering algorithm and thus an unsupervised learning method. Using this technique needs no explicit examples. Unclassified examples are parted in a hierarchy of natural groups by this procedure. This approach was also tested with synthetic data and reveals the following characterization of a landmark. Since the algorithm subdivides the records into similar groups, an instance with strongly different attribute values is separated from the others at a very high level in the decision tree. Because of its singularity, it is all alone in its group.

The methods are only tested with synthetic data, but the results were promising. Both lead to the desired result of identifying locally salient objects, that are likely to be distinguishable by (a set of) simple attributes. Details of the investigation are presented in (Elias, 2003).

6 VISIBILITY ANALYSIS USING LASER SCANNING DATASETS

6.1 Visibility Analysis

As noted in section 3.1, landmarks can only be of use for navigation purposes if they are sufficiently visible during the actual navigation process. Although some conclusions on visibility can be drawn from two-dimensional maps, important situations cannot be handled adequately. For example, figure 4(a) shows this case where the visibility of the tower on the right is not revealed.

The optimal case is of course when a full three-dimensional city model is available. Then, visibility can be computed exactly, yielding even information on the visibility of single building faces. Buildings standing out behind other buildings are correctly identified (Fig. 4(b)). If the model is only 2.5D, consisting of buildings with a single height (i.e., flat roofs), this can already lead to quite incorrect results. For example, imagine a large, flat building containing a tower of small footprint widely visible. In this case, depending on the choice for the (single) building height of the 2.5D model, the visibility is rated either much too conservative or too optimistic.




Figure 4: Visibility analysis for buildings (gray boxes) standing along a street (black lines). The visibility cone is shown in dark gray. (a) Based on 2D groundplans. (b) Based on true 3D geometry. (c) Based on a discrete DSM.

However, we can do better if we base the visibility analysis directly on the DSM from laser scanning. We will not obtain "beautiful" visualizations but instead a rather good estimate on which buildings can be seen from any viewpoint (Fig. 4(c)). We realized this approach as follows. For any viewpoint, the position and viewing direction define the exterior orientation of a virtual camera of given horizontal and vertical viewing angle. This virtual camera represents the driver's view. The height is derived from the DSM itself, whereas the viewing angle can be obtained from the orientation of the corresponding street segment in the GDF dataset.

The virtual image plane is then rastered, each pixel defining a ray in object space. All the rays are traced in object space to determine intersections with the DSM. For each hit, the corresponding object number is obtained by a lookup in an image containing rastered ground plan id's. Although this method is similar to "ray tracing" used in computer graphics and often assumed to be computationally expensive, it is actually quite fast since (a) we are interested only in the first hit of the ray, and (b) the DSM is 2.5D only, so each column in the virtual image plane can be computed efficiently from bottom to top, marching in increasing distance in object space. Figure 5 shows some examples. Clearly, one can see how larger objects such as towers can be identified sticking out behind other buildings.

6.2 Tracking Visibility

In the last section, visibility was computed for a single view. However, landmarks selected for a routing instruction must be visible during the entire manoeuvre. This can be checked by tracking the visibility of objects along the trajectory defined by the corresponding manoeuvre. For our first experiment, we use only a crude approximation for the visibility, namely the area covered by the projection of the corresponding object on the virtual image plane.

Figure 6 shows an example. We assume that the white polygon is the trajectory we want the driver to use. The question then is if the town hall, identified to be a landmark by the methods of section 5, is a suitable object which can be used in a landmark-based instruction such as 'pass to the right of the town hall'. To this end, our algorithm traces the entire trajectory, generating virtual views at equidistantly spaced positions and in the orientation defined by the trajectory. For each such view, the area covered by each object on the virtual image plane is determined.




Figure 6: Example trajectory, top view.

Figure 7 shows a plot of all those areas along the trajectory of figure 6. One can see the typical 'peaked' curves generated as objects appear, grow larger and finally disappear as the viewing position passes by. In this special case, one sees also that many objects become visible around frame number 65, which is when the view widens as the position leaves the narrow street and enters the plaza in front of the town hall.

In order to answer if the town hall is a suitable object, a look on figure 7 reveals that the corresponding curve (shown in bold red) is largest for frame numbers 65 to 115 (with a small exception around frame 100), i.e. the town hall is the largest object in the driver's view. Moreover, the curve is larger than zero starting from frame number 13, which means that the town hall is – at least partly – visible about 100 meters ahead of the position where the plaza is entered (which could be a decision point). Thus, in this case we can verify both that the object appears large and that it appears early enough in the driver's view.

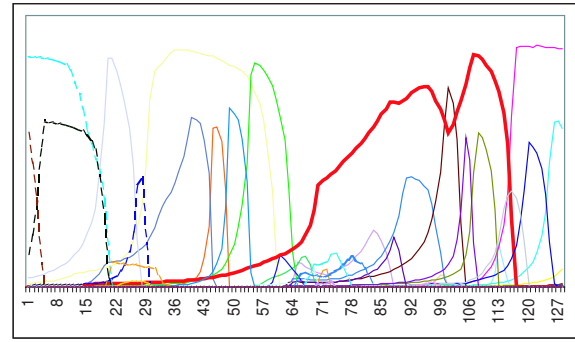


Figure 7: Visibility plotted over frame number. The frames were taken at equidistant 2 m intervals along the defined trajectory from figure 6. The visibility of the town hall is shown in bold red.

7 CONCLUSION AND OUTLOOK

In this paper, we have outlined how landmarks can be extracted and evaluated using existing GIS and laser scanning data. As for the extraction, we have investigated two different methods based on data mining to reveal prominent buildings. In order to evaluate the usefulness for navigation instructions, we used a visibility analysis based on DSM data from laser scanning.

Both data mining procedures have still to be tested with real data sets. The results will verify if they lead to appropriate landmarks in the real world. In addition, the analysis process has to be extended to different object types (traffic constructions, parks, sporting facilities, etc.) extracted for example from ATKIS data. Methods for data preprocessing of different object types and categories, and problems when different data mining algorithms are applied to the same data set, have to be investigated. The reliability of the extracted landmarks has to be determined by a quality measure to avoid ambiguous landmarks misleading the user.

More route-dependent aspects to determine real landmarks have to be investigated: The influence of the users moving direction and visibility on the quality of landmarks. As we only used the "virtual image size" to rate an object's visibility, there is much room for improvement. For example, from the virtual image, one can also obtain information on the distance, if the object is sticking out behind another, closer object, and if it is part of the silhouette. First pulse laser scan measurements could be integrated to get a better approximation for the occlusion caused by trees. The DSM could also be used to feed additional information to the extraction, for example, small towers sticking out behind a larger building could be identified. The implementation of the visibility tracking could also use equidistant time sampling instead of space sampling, based on assumed vehicle speeds in the vicinity of intersections. Finally, it would be interesting to investigate to what extent POI's already existent in GDF datasets could be used for the visibility analysis.

ACKNOWLEDGEMENTS

We thank the Stadtmessungsamt Stuttgart (surveying and mapping authority of the city of Stuttgart) for providing the cadastral map.

REFERENCES

AdV, 2003. Authoritative topographic cartographic information system (atkis). <http://www.adv-online.de/english/products/atkis.htm>. Last Visit: 02.04.03.




Figure 5: Some examples of the visibility computation. Left: top view, showing the DSM and the cone of visibility. Right: virtual image where each color corresponds to an object (building) number. (a) Market place of Stuttgart, surrounded by buildings. On the left, the town hall with tower. In the center background, the tower of the 'Stiftskirche' (church). (b) View from the 'Schlossplatz' (Stuttgart). On the left, a view along the 'Königstraße' with the tower of the main station in the background. The large building in the center is the new palace.

Baltsavias, E., Grün, A., and VanGool, L. (eds), 2001. Automatic Extraction of Man-Made Objects from Aerial and Satellite Images (III). Balkema Publishers, Netherlands.

Baltsavias, E. P., 1999. Airborne laser scanning: existing systems and firms and other resources. *ISPRS Journal of Photogrammetry and Remote Sensing* 54, pp. 164–198.

Baltsavias, E. P., Wehr, A. and Lohr, U. (eds), 1999. *ISPRS Journal of Photogrammetry and Remote Sensing, Special Issue on Airborne Laser Scanning*. Vol. 54, Elsevier.

Burnett, G. E., 1998. "Turn Right at the King's Head" Drivers' Requirements for Route Guidance Information. PhD thesis, Loughborough University, UK.

Burnett, G., Smith, D. and May, A., 2001. Supporting the navigation task: Characteristics of 'good' landmarks. In: M. A. Hanson (ed.), *Contemporary Ergonomics*, Taylor & Francis, pp. 441–446.

Elias, B., 2002. Automatic derivation of location maps. In: *GeoSpatial Theory, Processing and Applications*, Vol. 34/4, Ottawa, Canada.

Elias, B., 2003. Extracting landmarks with data mining methods. Technical report, ikg - Institute of Cartography and Geoinformatics, University of Hannover.

Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P. and Uthurusamy, R. (eds), 1996. *Advances in Knowledge Discovery and Data Mining*. AAAI Press/ The MIT Press, Menlo Park, Californien.

Frawley, W. J., Piatetsky-Shapiro, G. and Matheus, C., 1991. *Knowledge Discovery in Databases*. AAAI Press/ The MIT Press, Menlo Park, Californien, chapter Knowledge Discovery in Databases: An Overview, pp. 1–27.

Geographic Data Files 3.0, 1995. Technical report, European Committee for Standardization, CEN TC 278.

Golledge, R. D., 1999. Wayfinding Behavior. John Hopkins Press, chapter Human Wayfinding and Cognitive Maps, pp. 5–45.

Grün, A., Baltsavias, E. and Henricsson, O. (eds), 1997. Automatic Extraction of Man-Made Objects from Aerial and Space Images (II). Birkhäuser, Basel.

Grün, A., Kübler, O. and Agouris, P. (eds), 1995. Automatic Extraction of Man-Made Objects from Aerial and Space Images. Birkhäuser, Basel.

Kraus, K. and Rieger, W., 1999. Processing of laser scanning data for wooded areas. In: D. Fritsch and R. Spiller (eds), *Photogrammetric Week 99*, Wichmann Verlag, pp. 221–231.

Lovelace, K., Hegarty, M. and Montello, D., 1999. Elements of good route directions in familiar and unfamiliar environments. In: C. Freksa and D. Mark (eds), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science*, Springer Verlag, pp. 65–82.

Michon, P.-E. and Denis, M., 2001. When and why are visual landmarks used in giving directions? In: D. Montello (ed.), *COSIT 2001*, Springer Verlag, pp. 292–305.

Quinlan, J. R., 1986. Induction of decision trees. *Machine Learning*.

Raubal, M. and Winter, S., 2002. Enriching wayfinding instructions with local landmarks. In: M. Egenhofer and D. Mark (eds), *Geographic Information Science, Lecture Notes in Computer Science*, Vol. 2478, Springer Verlag, pp. 243–259.

Sorrows, M. and Hirtle, S., 1999. The nature of landmarks for real and electronic spaces. In: C. Freksa and D. Mark (eds), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science*, Springer Verlag, pp. 37–50.

Tversky, B. and Lee, P., 1999. Pictorial and verbal tools for conveying routes. In: C. Freksa and D. Mark (eds), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science*, Springer Verlag, pp. 51–64.

Winter, S., 2002. Ontologisches modellieren von routen für mobile navigationsdienste. In: F. Kelnhofer and M. Lechthaler (eds), *Telekartographie und Location Based Services, Geowissenschaftliche Mitteilungen*, Vol. 58, Schriftenreihe der Studienrichtung Vermessungswesen und Geoinformation, TU Wien, pp. 111–124.

Witten, I. H. and Eibe, F., 1999. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann, San Francisco.

Zhao, Y., 1997. *Vehicle Location and Navigation Systems*. Artech House, Inc. Boston, London.